

# The View-Update Problem for Indefinite Databases

Luciano Caroprese<sup>1</sup>, Irina Trubitsyna<sup>1</sup>, Mirosław Truszczyński<sup>2</sup>, and Ester Zumpano<sup>1</sup>

<sup>1</sup> DEIS, Università della Calabria, 87030 Rende, Italy  
caroprese|irina|zumpano@deis.unical.it

<sup>2</sup> Department of Computer Science, University of Kentucky, Lexington, KY 40506, USA  
mirek@cs.uky.edu

**Abstract.** This paper introduces and studies a declarative framework for updating views over *indefinite databases*. An indefinite database is a database with null values that are represented, following the standard database approach, by a single *null* constant. The paper formalizes views over such databases as indefinite *deductive* databases, and defines for them several classes of database repairs that realize view-update requests. Most notable is the class of *constrained repairs*. Constrained repairs change the database “minimally” and avoid making arbitrary commitments. They narrow down the space of alternative ways to fulfill the view-update request to those that are grounded, in a certain strong sense, in the database, the view and the view-update request.

## 1 Introduction

A typical database system is large and complex. Users and applications rarely can access the entire system directly. Instead, it is more common that access is granted in terms of a *view*, a *virtual* database consisting of relations defined by a query to the *stored and maintained* database. Querying a view does not present a conceptual problem. In contrast, another key task, *view updating*, poses major challenges.

*Example 1.* Let  $D = \{q(a,b)\}$  be a database over relation symbols  $q$  and  $r$ , where the relation  $r$  has arity three and is currently empty. Let us consider the view over  $D$  given by the Datalog program  $P = \{p(X) \leftarrow q(X,Y), r(X,Y,Z)\}$ . That view consists of a single unary relation  $p$ . Given the present state of  $D$ , the view is empty.

To satisfy the request that  $p(a)$  holds in the view (as it is now, it does not), one needs to update the database  $D$ . Such update consists of executing update actions that specify facts to insert to and to delete from  $D$ . These update actions (in a simplified setting that we consider for now) are “signed” facts  $+F$  and  $-G$ , where  $+F$  stands for “insert  $F$ ” and  $-G$  stands for “delete  $G$ .” In our case, the set of update actions  $\{-q(a,b), +q(a,a), +r(a,a,a)\}$  is a correct update to  $D$ . Executing it on  $D$  results in the database  $D' = \{q(a,a), r(a,a,a)\}$ , which has the desired property that  $p(a)$  holds in the view determined by  $P$ . There are also other ways to satisfy the user’s request, for instance:  $\{+r(a,b,a)\}$  and  $\{+q(c,d), +r(c,d,d)\}$ , where  $c$  and  $d$  are any elements of the domain of the database.  $\square$

As this example suggests, view updating consists of translating a *view-update request*, that is, an update request against the view, into an *update*, a set of update actions

against the stored (extensional) database. The example highlights the basic problem of view updating. It may be (in fact, it is common), that a view-update request can be fulfilled by any of a large number of database updates. One of them has to be committed to and executed. Thus, developing methods to automate the selection process or, at least, aid the user in making the choice is essential and has been one of the central problems of view updating [9, 2, 13, 11]. That problem is also the focus of our paper.<sup>1</sup>

To restrict the space of possible database updates to select from, it is common to consider only those that accomplish the view-update request and are in some sense minimal. That reduces the space of updates. For instance, the update  $\{-q(a, b), +q(a, a), +r(a, a, a)\}$  in Example 1 is not minimal. We can remove  $-q(a, b)$  from it and what remains is still an update that once executed ensures that  $p(a)$  holds in the view. Minimal updates fulfilling the view-update request are commonly called *repairs*.

Even imposing minimality may still leave a large number of candidate repairs. In Example 1, updates  $\{+r(a, b, \psi)\}$ , where  $\psi$  is any domain element, are repairs, and there are still more. As long as we insist on the completeness of the repaired database, there is little one can do at this point but ask the user to *select* one.

The situation changes if we are willing to allow indefiniteness (incomplete information) in databases. Given the “regular” structure of the family of repairs above, one could represent it by a single *indefinite* repair,  $\{+r(a, b, \psi)\}$  regarding  $\psi$  as a distinct *null* value standing for some unspecified domain elements. The choice facing the user substantially simplifies as possibly infinitely many repairs is reduced to only one.

This approach was studied by Farré et al. [5], where the terminology of Skolem constants rather than null values was used. However, while seemingly offering a plausible solution to the problem of multiple repairs, the approach suffers from two drawbacks. First, the use of multiple Skolem constants (essentially, multiple null values) violates the SQL standard [15]. Second, the approach assumes that the initial database is complete (does not contain null values). Thus, while the approach might be applied once, iterating it fails as soon as an indefinite database is produced, since no guidance on how to proceed in such case is given.

We consider the view-update problem in the setting when *both* the original and the updated databases are indefinite (may contain occurrences of null values). To be compliant with the SQL standard, we allow a single null value only. We denote it by  $\perp$  and interpret it as expressing the existence of unknown values for each attribute it is used for [10]. Typically, a database is represented as a set of facts. We propose to represent indefinite databases by *two* sets of facts that we interpret by means of a two-level closed-world assumption tailored to the case of indefinite information. We interpret indefinite databases in terms of their *possible worlds*. We extend the possible-world semantics to the setting of views, which we formalize in terms of *indefinite deductive databases*, and apply the resulting formalism to state and study the view-updating problem.

We then turn attention to the problem of multiple repairs. In general, using null values to encode multiple repairs is still not enough to eliminate them entirely (cf. our

---

<sup>1</sup> In some cases no update satisfies the view-update request. The question whether the lack of an appropriate update is caused by errors in the design of the view, in the extensional database, or in the view-update request is interesting and deserves attention, but is outside the scope of the present work.

discussion above), and some level of the user’s involvement may be necessary. Therefore, it is important to identify principled ways to narrow down the space of repairs for the user to consider. As the setting we consider here, when both the original and the updated databases may contain null values has not been considered before, we propose a concept of minimality tailored precisely to that situation. The primary concern is to minimize the set of new constants introduced by an update. The secondary concern is to minimize the degree of the semantic change. The resulting notion of minimality yields the notion of a *repair*.

Our concept of minimality leads us also to the concept of a *relevant* repair, an update that introduces no new constants and minimizes the degree of change. In Example 1,  $\{+r(a, b, \perp)\}$ ,  $\{+r(a, a, \perp), +q(a, a)\}$  and  $\{+r(a, c, \perp), +q(a, c)\}$ , where  $c$  is an element of the database domain other than  $a$  and  $b$ , are all repairs. The first two are obviously relevant, the third one is not.

Some occurrences of non-nullary constants in a relevant repair may still be “ungrounded” or “arbitrary,” that is, replacing them with another constant results in a repair. For instance, replacing in the relevant repair  $\{+r(a, a, \perp), +q(a, a)\}$  the second and the forth occurrences of  $a$  with a fresh constant  $c$  yields  $\{+r(a, c, \perp), +q(a, c)\}$ , an update that is a repair. Intuitively, “arbitrary” occurrences of constants, being replaceable, are not forced by the information present in the view (deductive database) and in the view-update request. By restricting relevant repairs to those *without* arbitrary occurrences of constants we arrive at the class of *constrained* repairs. In the view-update problem considered in Example 1, there is only *one* constrained repair,  $\{+r(a, b, \perp)\}$ .

Finally, we study the complexity of the problems of the existence of repairs, relevant repairs and constrained repairs. We obtain precise results for the first two classes and an upper bound on the complexity for the last class.

To summarize, our main contributions are as follows. We propose a two-set representation of indefinite database that is more expressive than the standard one. We define the semantics and the operation of updating indefinite databases (Section 2). We define views over indefinite databases (indefinite deductive databases), and generalize the semantics of indefinite databases to views (Section 3). We state and study the view-update problem in the general setting when the initial and the repaired databases are indefinite. We propose a notion of minimality of an update and use it to define the concept of a repair. We address the problem of multiple repairs by defining relevant and constrained repairs (Section 4). We study the complexity of problems of existence of repairs, relevant repairs and constrained repairs (Section 5). The proofs are in the appendix.

## 2 Indefinite Databases

We consider a finite set  $\Pi$  of relation symbols and a set  $Dom$  of constants that includes a designated element  $\perp$ , called the *null value*. We define  $Dom_d = Dom \setminus \{\perp\}$ . Normally, we assume that  $Dom$  is an infinite countable set. However, for the sake of simplicity, in several of the examples the set  $Dom$  is finite.

Some predicates in  $\Pi$  are designated as *base* (or *extensional*) predicates and all the remaining ones are understood as *derived* (or *intensional*) predicates. A *term* is a constant from  $Dom$  or a variable. An *atom* is an expression of the form  $p(t_1, \dots, t_k)$ , where  $p \in \Pi$  is a *predicate symbol* of arity  $k$  and  $t_i$ ’s are terms. An atom is *ground* if

it does not contain variables. We refer to ground atoms as *facts*. We denote the set of all facts by  $At$ . We call facts defined in terms of base and derived predicates *base facts* and *derived facts* respectively. A fact is *definite* if it does not contain occurrences of  $\perp$ . Otherwise, it is *indefinite*. Given a set  $S$  of atoms, we define  $Dom(S)$  (resp.  $Dom_d(S)$ ) as the set of constants in  $Dom$  (resp.  $Dom_d$ ) occurring in  $S$ . For every two tuples of terms  $t = (t_1, \dots, t_k)$  and  $t' = (t'_1, \dots, t'_k)$  and every  $k$ -ary predicate symbol  $p \in \Pi$ , we write  $t \preceq t'$  and  $p(t) \preceq p(t')$  if for every  $i$ ,  $1 \leq i \leq k$ ,  $t_i = t'_i$  or  $t_i = \perp$ . We say in such case that  $t'$  and  $p(t')$  are *at least as informative* as  $t$  and  $p(t)$ , respectively. If, in addition,  $t \neq t'$ , we write  $t \prec t'$  and  $p(t) \prec p(t')$ , and say that  $t'$  and  $p(t')$  are *more informative* than  $t$  and  $p(t)$ . Sometimes, we say “at most as informative” and “less informative,” with the obvious understanding of the intended meaning. We also define  $t$  and  $t'$  (respectively,  $p(t)$  and  $p(t')$ ) to be *compatible*, denoted by  $t \approx t'$  (respectively,  $p(t) \approx p(t')$ ), if for some  $k$ -tuple  $s$  of terms,  $t \preceq s$  and  $t' \preceq s$ . Finally, for a set  $D \subseteq At$ , we define

$$\begin{aligned} D^\Downarrow &= \{a \mid \text{there is } b \in D \text{ s.t. } a \preceq b\} & D^\Uparrow &= \{a \mid \text{there is } b \in D \text{ s.t. } b \preceq a\} \\ D^\approx &= \{a \mid \text{there is } b \in D \text{ s.t. } b \approx a\} & D^\sim &= D^\approx \setminus D^\Downarrow. \end{aligned}$$

To illustrate, let  $q$  be a binary relation symbol and  $Dom = \{\perp, 1, 2\}$ . Then:

$$\begin{aligned} \{q(1, \perp)\}^\Downarrow &= \{q(1, \perp), q(\perp, \perp)\} & \{q(1, \perp)\}^\Uparrow &= \{q(1, \perp), q(1, 1), q(1, 2)\} \\ \{q(1, \perp)\}^\approx &= \{q(1, \perp), q(1, 1), q(1, 2), q(\perp, 1), q(\perp, 2), q(\perp, \perp)\} \\ \{q(1, \perp)\}^\sim &= \{q(1, 1), q(1, 2), q(\perp, 1), q(\perp, 2)\}. \end{aligned}$$

We note also note that  $D^\approx = (D^\Uparrow)^\Downarrow$ .

In the most common case, databases are finite subsets of  $At$  that contain *definite* facts only. The semantics of such databases is given by the *closed-world assumption* or CWA [12]: a definite fact  $q$  is *true* in a database  $D$  if  $q \in D$ . Otherwise,  $q$  is *false* in  $D$ . We are interested in databases that may contain indefinite facts, too. Generalizing, we will *for now* assume that an indefinite database is a finite set of possibly indefinite atoms. The key question is that of the semantics of indefinite databases.

Let  $D$  be an indefinite database. Clearly, all facts in  $D$  are true in  $D$ . In addition, any fact that is less informative than a fact in  $D$  is also true in  $D$ . Indeed, each such fact represents an existential statement, whose truth is established by the presence of a more informative fact in  $D$  (for instance, the meaning of  $p(\perp)$  is that there is an element  $c$  in the domain of the database such that  $p(c)$  holds; if  $p(1) \in D$ , that statement is true). Summarizing, every fact in  $D^\Downarrow$  is true in  $D$ .

By CWA adapted for the setting of indefinite databases [10], facts that are not in  $D^\Downarrow$  are *not* true in  $D$ , as  $D$  contains no evidence to support their truth. Those facts among them that are compatible with facts in  $D$  (in our notation, facts in  $D^\sim$ ), might actually be true, but the database just does not know that. Of course, they may also be false, the database does not exclude that possibility either. They are regarded as *unknown*. By CWA again, the facts that are not compatible with any fact in  $D$  are false in  $D$ , as  $D$  provides no explicit evidence otherwise.

The simple notion of an indefinite database, while intuitive and having a clear semantics, has a drawback. It has a limited expressive power. For instance, there is no database  $D$  to represent our knowledge that  $p(1)$  is false and that there is some definite  $c$  such that  $p(c)$  holds (clearly, this  $c$  is not 1). To handle such cases, we introduce a more general concept of an indefinite database, still using CWA to specify its meaning.

**Definition 1.** An indefinite database (a database, for short) is a pair  $\mathcal{J} = \langle D, E \rangle$ , where  $D$  and  $E$  are finite sets of (possibly indefinite) facts.  $\square$

The intended role of  $D$  is to represent all facts that are true in the database  $\langle D, E \rangle$ , while  $E$  is meant to represent *exceptions*, those facts that normally would be unknown, but are in fact false (and the database knows it). More formally, the semantics of indefinite databases is presented in the following definition.

**Definition 2.** Let  $\langle D, E \rangle$  be a database and let  $q \in At$  be a fact. Then: (1)  $q$  is true in  $\langle D, E \rangle$ , written  $\langle D, E \rangle \models q$ , if  $q \in D^\downarrow$ ; (2)  $q$  is unknown in  $\langle D, E \rangle$ , if  $q \in (D^\sim \setminus D^\downarrow) \setminus E^\uparrow$  ( $= D^\sim \setminus E^\uparrow$ ); (3)  $q$  is false in  $\langle D, E \rangle$ , denoted  $\langle D, E \rangle \models \neg q$ , in all other cases, that is, if  $q \notin D^\sim$  or if  $q \in D^\sim \cap E^\uparrow$ .  $\square$

The use of  $E^\uparrow$  in the definition (items (2) and (3)) reflects the property that if an atom  $a$  is false then every atom  $b$  at least as informative as  $a$  must be false, too.

We denote the sets of all facts that are true, unknown and false in a database  $\mathcal{J} = \langle D, E \rangle$  by  $\mathcal{J}_t$ ,  $\mathcal{J}_u$  and  $\mathcal{J}_f$ , respectively. Restating the definition we have:

$$\mathcal{J}_t = D^\downarrow, \quad \mathcal{J}_u = D^\sim \setminus E^\uparrow, \quad \text{and} \quad \mathcal{J}_f = (At \setminus D^\sim) \cup (D^\sim \cap E^\uparrow).$$

*Example 2.* The knowledge that  $p(c)$  holds for some constant  $c$  and that  $p(1)$  is false can be captured by the database  $\langle \{p(\perp)\}, \{p(1)\} \rangle$ . The database  $\langle \{q(\perp, \perp), q(1, 1)\}, \{q(1, \perp)\} \rangle$  specifies that the atoms  $q(1, 1)$ ,  $q(1, \perp)$ ,  $q(\perp, 1)$  and  $q(\perp, \perp)$  are true, that all definite atoms  $q(1, d)$ , with  $d \neq 1$ , are false, and that all other atoms  $q(c, d)$  are unknown. While the fact that  $q(\perp, \perp)$  is true follows from the fact that  $q(1, 1)$  is true, the presence of the former in the database is not redundant, that is, the database without it would have a different meaning. Namely,  $q(\perp, \perp)$  makes all atoms  $q(a, b)$  potentially unknown (with some of them true or false due to other elements of the database).  $\square$

**Definition 3.** A set  $W$  of definite facts is a possible world for a database  $\mathcal{J} = \langle D, E \rangle$  if  $\mathcal{J}_t \subseteq W^\downarrow$  ( $W$  “explains” all facts that are true in  $\mathcal{J}$ ),  $W \subseteq \mathcal{J}_t \cup \mathcal{J}_u$  (only definite facts that are true or unknown appear in a possible world).  $\square$

Databases represent sets of possible worlds. For a database  $\mathcal{J}$ , we write  $\mathcal{W}(\mathcal{J})$  to denote the family of all possible worlds for  $\mathcal{J}$ . Due to the absence of indefinite facts in  $W \in \mathcal{W}(\mathcal{J})$ , every fact in  $At$  is either true (if it belongs to  $W^\downarrow$ ) or false (otherwise) w.r.t.  $W$ . Extending the notation we introduced earlier, for a possible world  $W$  and  $a \in At$  we write  $W \models a$  if  $a \in W^\downarrow$  and  $W \models \neg a$ , otherwise. The following proposition shows that the semantics of a database can be stated in terms of its possible worlds.

**Proposition 1.** Let  $\mathcal{J}$  be a database and  $q$  a fact. Then  $q \in \mathcal{J}_t$  if and only if  $W \models q$ , for every  $W \in \mathcal{W}(\mathcal{J})$ , and  $q \in \mathcal{J}_f$  if and only if  $W \models \neg q$ , for every  $W \in \mathcal{W}(\mathcal{J})$ .  $\square$

Updating a database  $\langle D, E \rangle$  consists of executing on it *update actions*: inserting a base fact  $a$  into  $D$  or  $E$ , and deleting a base fact  $a$  from  $D$  or  $E$ . We denote them by  $+a^D$ ,  $+a^E$ ,  $-a^D$  and  $-a^E$ , respectively. For a set  $U$  of update actions, we define  $U_+^D = \{a \mid +a^D \in U\}$ ,  $U_+^E = \{a \mid +a^E \in U\}$ ,  $U_-^D = \{a \mid -a^D \in U\}$ , and  $U_-^E = \{a \mid -a^E \in U\}$ . To be executable, a set  $U$  of update actions must not contain *contradictory* update

actions:  $+a^D$  and  $-a^D$ , or  $+a^E$  and  $-a^E$ . A contradiction-free set  $U$  of update actions is an *update*. We denote the set of all updates (in the fixed language we consider) by  $\mathcal{W}$ .

We now define the operation to update a database, that is, to apply an update to it.

**Definition 4.** Let  $\mathcal{J} = \langle D, E \rangle$  be an indefinite database and  $U$  an update. We define  $\mathcal{J} \circ U$  as the database  $\langle D', E' \rangle$ , where  $D' = (D \cup U_+^D) \setminus U_-^D$  and  $E' = (E \cup U_+^E) \setminus U_-^E$ .  $\square$

### 3 Indefinite Deductive Databases

*Integrity constraints* (ICs, for short) are first-order sentences in the language  $\mathcal{L}$  determined by the set of predicates  $\Pi$  and by the set  $Dom_d$  of definite constants. A database with integrity constraints is a pair  $\langle \mathcal{J}, \eta \rangle$ , where  $\mathcal{J}$  is a database and  $\eta$  is a set of ICs. Possible worlds can be regarded as interpretations of the language  $\mathcal{L}$  with  $Dom_d$  as their domain. This observation and Proposition 1 suggest a definition of the semantics of databases with ICs.

**Definition 5.** Let  $\langle \mathcal{J}, \eta \rangle$  be a database with ICs. A possible world  $W \in \mathcal{W}(\mathcal{J})$  is a possible world for a database  $\langle \mathcal{J}, \eta \rangle$  if  $W$  satisfies every integrity constraint in  $\eta$ . We denote the set of possible worlds of  $\langle \mathcal{J}, \eta \rangle$  by  $\mathcal{W}(\mathcal{J}, \eta)$ . A database with ICs,  $\langle \mathcal{J}, \eta \rangle$ , is consistent if  $\mathcal{W}(\mathcal{J}, \eta) \neq \emptyset$ . Otherwise,  $\langle \mathcal{J}, \eta \rangle$  is inconsistent.

A fact  $q$  is true in  $\langle \mathcal{J}, \eta \rangle$  if  $W \models q$ , for every  $W \in \mathcal{W}(\mathcal{J}, \eta)$ ;  $q$  is false in  $\langle \mathcal{J}, \eta \rangle$  if  $W \models \neg q$ , for every  $W \in \mathcal{W}(\mathcal{J}, \eta)$ ; otherwise,  $q$  is unknown in  $\langle \mathcal{J}, \eta \rangle$ .  $\square$

*Example 3.* Let us consider the database  $\mathcal{J} = \langle \{p(1), p(2), q(\perp)\}, \emptyset \rangle$  (there are no exceptions) and let  $\eta = \{\forall X(q(X) \rightarrow p(X))\}$ . Possible worlds of  $\mathcal{J}$  include  $\{p(1), p(2), q(1)\}$ ,  $\{p(1), p(2), q(2)\}$  and  $\{p(1), p(2), q(3)\}$ . The first two satisfy the integrity constraint, the third one does not. Thus, only the first two are possible worlds of  $\langle \mathcal{J}, \eta \rangle$ . Since  $p(1)$  belongs to all possible worlds of  $\langle \mathcal{J}, \eta \rangle$ ,  $p(1)$  is true in  $\langle \mathcal{J}, \eta \rangle$ . Further,  $p(3)$  is false in every possible world of  $\mathcal{J}$  and so also in every possible world of  $\langle \mathcal{J}, \eta \rangle$ . Thus,  $p(3)$  is false in  $\langle \mathcal{J}, \eta \rangle$ . Lastly, we note that  $q(1)$  and  $q(2)$  are unknown in  $\langle \mathcal{J}, \eta \rangle$ , while  $q(3)$  is false (due to  $p(3)$  being false and the integrity constraint).  $\square$

We note that the possible-world semantics can capture additional information contained in integrity constraints. In Example 3, the semantics derives that  $q(3)$  is false in  $\langle \mathcal{J}, \eta \rangle$  even though this knowledge is not present in the database  $\mathcal{J}$ .

The concepts of an update and of the operation to execute an update on a database extend literally to the case of databases with ICs.

Following Ullman [16], *views* are safe Datalog<sup>+</sup> programs. We use the standard terminology and talk about (Datalog<sup>+</sup>) *rules*, and *bodies* and *heads* of rules. A rule is *safe* if each variable occurring in the head or in a negative literal in the body also occurs in a positive literal in the body. A Datalog<sup>+</sup> program is safe if each rule is *safe*. We assume that views do not contain occurrences of  $\perp$ . The semantics of Datalog<sup>+</sup> programs is given in terms of *answer sets* [6, 7]. A precise definition of that semantics is immaterial to our study and so we do not provide the details.

**Definition 6.** An indefinite deductive database (from now, simply, a *deductive database*) is a tuple  $\mathcal{D} = \langle \mathcal{J}, \eta, P \rangle$ , where  $\mathcal{J}$  is a database,  $\eta$  is a set of integrity constraints, and



$P$  is a safe Datalog<sup>-</sup> program (the specification of a view) such that no predicate occurring in the head of a rule in  $P$  is a base predicate.  $\square$

Clearly, a deductive database with the empty view is a database with ICs, and a deductive database with the empty view and no ICs is simply a database.

**Definition 7.** A deductive database  $\mathcal{D} = \langle \mathcal{I}, \eta, P \rangle$  is consistent if  $\langle \mathcal{I}, \eta \rangle$  is consistent and for every possible world  $W \in \mathcal{W}(\mathcal{I}, \eta)$ , the program  $W \cup P$  has answer sets. We denote the family of all those answer sets by  $\mathcal{W}(\mathcal{D})$  or  $\mathcal{W}(\mathcal{I}, \eta, P)$ . We call elements of  $\mathcal{W}(\mathcal{D})$  possible worlds of  $\mathcal{D}$ .  $\square$

There is an alternative to our concept of consistency. One could define a deductive database  $\langle \mathcal{I}, \eta, P \rangle$  as consistent if for *at least one* world  $W \in \mathcal{W}(\mathcal{I}, \eta)$ , the program  $W \cup P$  has an answer set. That concept of consistency would allow situations where for some possible worlds of  $\langle \mathcal{I}, \eta \rangle$ , one of which could be a description of the real world, the view  $P$  does not generate any meaningful virtual database. Our concept of consistency is more robust. It guarantees that the user can have a view of a database no matter how the real world looks like, that is, which of the possible worlds describes it.

*Example 4.* Let  $\mathcal{D} = \langle \mathcal{I}, \eta, P \rangle$  be a deductive database, where  $\mathcal{I} = \langle \{p(\perp)\}, \emptyset \rangle$ ,  $\eta = \emptyset$  and  $P = \{t \leftarrow p(1), p(2), \text{not } t\}$ , for some derived ground atom  $t$ . Every non-empty set  $W \subseteq \{p(u) \mid u \in \text{Dom}_d\}$  is a possible world of  $\langle \mathcal{I}, \eta \rangle$ . In particular, the set  $\{p(1), p(2)\}$  is a possible world of  $\langle \mathcal{I}, \eta \rangle$ . Since the program  $P \cup \{p(1), p(2)\}$  has no answer sets,  $\mathcal{D}$  is inconsistent (according to our definition). If  $\mathcal{D}' = \langle \mathcal{I}, \{p(1) \wedge p(2) \rightarrow \perp\}, P \rangle$ , then the integrity constraint in  $\mathcal{D}'$  eliminates the offending possible world and one can check that for every possible world  $W$  of  $\langle \mathcal{I}, \{p(1) \wedge p(2) \rightarrow \perp\} \rangle$ ,  $P \cup W$  has an answer set. Thus,  $\mathcal{D}'$  is consistent.<sup>2</sup>  $\square$

The concept of an update extends in a natural way to deductive databases. If  $U$  is an update, and  $\mathcal{D} = \langle \mathcal{I}, \eta, P \rangle$ , we define the result of updating  $\mathcal{D}$  by  $U$  by  $\mathcal{D} \circ U = \langle \mathcal{I} \circ U, \eta, P \rangle$ .

Next, we define the semantics of a deductive database  $\mathcal{D} = \langle \mathcal{I}, \eta, P \rangle$ , again building on the characterization given by Proposition 1.

**Definition 8.** A fact  $a \in \text{At}$  is true in a deductive database  $\mathcal{D} = \langle \mathcal{I}, \eta, P \rangle$ , denoted by  $\mathcal{D} \models a$ , if for every possible world  $W \in \mathcal{W}(\mathcal{D})$ ,  $W \models a$ ;  $a$  is false in  $\mathcal{D}$ , denoted by  $\mathcal{D} \models \neg a$ , if for every  $W \in \mathcal{W}(\mathcal{D})$ ,  $W \models \neg a$ ;  $a$  is unknown in  $\mathcal{D}$ , otherwise. We denote the truth value of  $a$  in  $\mathcal{D}$  by  $v_{\mathcal{D}}(a)$ .  $\square$

*Example 5.* Let  $\text{Dom} = \{\perp, 1, 2, 3\}$  and  $\mathcal{D} = \langle \mathcal{I}, \eta, P \rangle$  be a deductive database, where  $\mathcal{I} = \langle \{p(\perp)\}, \emptyset \rangle$ ,  $\eta = \{p(2) \rightarrow \perp\}$  and  $P = \{q(X) \leftarrow p(X)\}$ .

We have  $\mathcal{W}(\mathcal{I}, \eta) = \{\{p(1)\}, \{p(3)\}, \{p(1), p(3)\}\}$ . Each of the possible worlds in  $\mathcal{W}(\mathcal{I}, \eta)$ , when extended with the view  $P$ , gives rise to a program that has answer sets. Thus,  $\mathcal{D}$  is consistent. Moreover, the possible worlds for  $\mathcal{D}$  are  $\{p(1), q(1)\}$ ,  $\{p(3), q(3)\}$  and  $\{p(1), q(1), p(3), q(3)\}$  (in this case, one for each possible world in  $\mathcal{W}(\mathcal{I}, \eta)$ ). It follows that  $p(\perp)$  and  $q(\perp)$  are true,  $p(2)$  and  $q(2)$  are false, and  $p(1), q(1), p(3)$  and  $q(3)$  are unknown in  $\mathcal{D}$ .  $\square$

<sup>2</sup> We point out that in the paper, we overload the notation  $\perp$ . We use it to denote both the single null value in the language and the falsity symbol in the first-order language used for integrity constraints. Since the meaning is always clear from the context, no ambiguity arises.

## 4 View Updating

In the *view update problem*, the user specifies a *request*, a list of facts the user learned (observed) to be true or false, and wants the stored database to be updated to reflect it.<sup>3</sup>

**Definition 9.** A request over a deductive database  $\mathcal{D}$  is a pair  $\mathcal{S} = (\mathcal{S}_t, \mathcal{S}_f)$ , where  $\mathcal{S}_t$  and  $\mathcal{S}_f$  are disjoint sets of facts requested to be true and false, respectively.  $\square$

To fulfill a request we need an update which, when executed, yields a database such that the view it determines satisfies the request. We call such updates *weak repairs*.

**Definition 10.** Let  $\mathcal{D} = \langle \mathcal{I}, \eta, P \rangle$  be a deductive database and  $\mathcal{S}$  a request. An update  $U$  for  $\mathcal{S}$  is a weak repair for  $(\mathcal{D}, \mathcal{S})$  if  $U$  fulfills  $\mathcal{S}$ , that is, if for every  $a \in \mathcal{S}_t$ ,  $v_{\mathcal{D} \circ \mathcal{U}}(a) = \text{true}$  and for every  $a \in \mathcal{S}_f$ ,  $v_{\mathcal{D} \circ \mathcal{U}}(a) = \text{false}$ .  $\square$

We are primarily interested in updates that do not drastically change the database. One condition of being “non-drastic” is not to introduce new predicate or constant symbols. That leads us to the notion of a relevant weak repair.

**Definition 11.** Let  $\mathcal{D} = \langle \mathcal{I}, \eta, P \rangle$  be a deductive database and  $\mathcal{S}$  a request. A constant is relevant with respect to  $\mathcal{D}$  and  $\mathcal{S}$  if it occurs in  $\mathcal{D}$ , or  $\mathcal{S}$ , or if it is  $\perp$ . A predicate is relevant with respect to  $\mathcal{D}$  and  $\mathcal{S}$  if it occurs in  $\mathcal{D}$  or in  $\mathcal{S}$ . A weak repair  $U$  for  $(\mathcal{D}, \mathcal{S})$  is relevant if every constant and predicate occurring in  $U$  is relevant.  $\square$

More generally, a weak repair is “non-drastic” if it minimizes the change it incurs [14]. There are two aspects to the minimality of change: (1) minimizing the set of new predicate symbols and constants introduced by an update to the database (in the extreme case, no new symbols must be introduced, and we used that requirement to define relevant weak repairs above); (2) minimizing the change in the truth values of facts with respect to the database. Following the *Ockham’s Razor* principle to avoid introducing new entities unless necessary, we take the minimality of the set of new symbols as a primary consideration. To define the resulting notion of change minimality, we assume that the truth values are ordered  $\text{false} \leq \text{unknown} \leq \text{true}$ . Further, for a deductive database  $\mathcal{D}$ , a request set  $\mathcal{S}$  and an update  $U \in \mathcal{U}$  we define  $NC(\mathcal{D}, \mathcal{S}, U)$  as the set of non nullary constants that occur in  $U$  and not in  $\mathcal{D}$  and  $\mathcal{S}$ .

**Definition 12.** Let  $\mathcal{D} = \langle \mathcal{I}, \eta \rangle$  be a database with integrity constraints. For updates  $V, U \in \mathcal{U}$ , we define  $U \sqsubseteq V$  if:  $NC(\mathcal{D}, \mathcal{S}, U) \subset NC(\mathcal{D}, \mathcal{S}, V)$ , or  $NC(\mathcal{D}, \mathcal{S}, U) = NC(\mathcal{D}, \mathcal{S}, V)$  and for every base atom  $a$

1. if  $v_{\mathcal{D}}(a) = \text{true}$ , then  $v_{\mathcal{D} \circ \mathcal{U}}(a) \geq v_{\mathcal{D} \circ \mathcal{V}}(a)$
2. if  $v_{\mathcal{D}}(a) = \text{false}$ , then  $v_{\mathcal{D} \circ \mathcal{V}}(a) \geq v_{\mathcal{D} \circ \mathcal{U}}(a)$
3. if  $v_{\mathcal{D}}(a) = \text{unknown}$ , then  $v_{\mathcal{D} \circ \mathcal{U}}(a) = \text{unknown}$  or  $v_{\mathcal{D} \circ \mathcal{V}}(a) = v_{\mathcal{D} \circ \mathcal{U}}(a)$ .

We also define  $U \sqsubset V$  if  $U \sqsubseteq V$  and  $V \not\sqsubseteq U$ .  $\square$

<sup>3</sup> We do not allow requests that facts be *unknown*. That is, we only allow definite requests. While there may be situations when all the user learns about the fact is that it is unknown, they seem to be rather rare. In a typical situation, the user will learn the truth or falsity of a fact.



We now define the classes of *repairs* and *relevant repairs* as subclasses of the respective classes of weak repairs consisting of their  $\sqsubseteq$ -minimal elements.

**Definition 13.** Let  $\mathcal{D} = \langle \mathcal{I}, \eta, P \rangle$  be a deductive database and  $\mathcal{S}$  a request. A (relevant) repair for  $(\mathcal{D}, \mathcal{S})$  is a  $\sqsubseteq$ -minimal (relevant) weak repair for  $(\mathcal{D}, \mathcal{S})$ .  $\square$

We note that the existence of (weak) repairs does not guarantee the existence of relevant (weak) repairs. The observation remains true even if the view is empty.

*Example 6.* Let  $\mathcal{D} = \langle \mathcal{I}, \eta, P \rangle$ , where  $\mathcal{I} = \langle \emptyset, \emptyset \rangle$ ,  $\eta = \emptyset$  and  $P = \{t \leftarrow p(x), q(x)\}$ . If the request is  $(\{t\}, \emptyset)$ , then each repair is of the form  $\{+p(i)^D, +q(i)^D\}$ , for some  $i \in \text{Dom}_d$ . None of them is relevant. (We note that  $\{+p(\perp)^D, +q(\perp)^D\}$  is not a (weak) repair. The database resulting from the update would admit possible worlds of the form  $\{p(i), q(j)\}$ , where  $i \neq j$ . Clearly, the corresponding possible world of the view over any such database does not contain  $t$  and so the update does not fulfill the request.)  $\square$

Some relevant constants are not “forced” by the database and the request, that is, can be replaced by other constants. If such constants are present in a relevant (weak) repair, this repair is *arbitrary*. Otherwise, it is *constrained*. A formal definition follows.

**Definition 14.** Let  $\mathcal{D} = \langle \mathcal{I}, \eta, P \rangle$  be a deductive database and  $\mathcal{S}$  a request. A relevant (weak) repair  $U$  for  $(\mathcal{D}, \mathcal{S})$  is constrained if there is no non-nullary constant  $a$  in  $U$  such that replacing some occurrences of  $a$  in  $U$  with a constant  $b \neq a$  ( $b$  might be  $\perp$ ), results in a weak repair for  $(\mathcal{D}, \mathcal{S})$ .  $\square$

*Example 7.* Let  $\mathcal{D} = \langle \mathcal{I}, \eta, P \rangle$ , where  $\mathcal{I} = \langle \{p(1), h(2)\}, \emptyset \rangle$ ,  $\eta = \emptyset$  and  $P = \{t \leftarrow p(X), q(X); s \leftarrow r(X)\}$ . Let us consider the request  $\mathcal{S} = (\{s, t\}, \emptyset)$ . The updates  $\mathcal{R}_i = \{+q(1)^D, +r(i)^D\}$ ,  $i \in \{\perp, 1, 2\}$ , and  $\mathcal{R}'_i = \{+q(2)^D, +p(2)^D, +r(i)^D\}$ ,  $i \in \{\perp, 1, 2\}$ , are relevant weak repairs. One of them,  $\mathcal{R}_\perp$ , is constrained. Indeed, replacing in  $\mathcal{R}_\perp$  the unique occurrence of a non-nullary constant (in this case, 1) with any other constant does not yield a weak repair. On the other hand,  $\mathcal{R}_i$ ,  $i \in \{1, 2\}$ , and  $\mathcal{R}'_i$ ,  $i \in \{\perp, 1, 2\}$ , are not constrained. Indeed, replacing with 3 the second occurrence of 1 in  $\mathcal{R}_1$ , or the occurrence of 2 in  $\mathcal{R}_2$ , or both occurrences of 2 in  $\mathcal{R}'_i$  in each case results in a weak repair. Also weak repairs  $\mathcal{R}_i = \{+q(1)^D, +r(i)^D\}$  and  $\mathcal{R}'_i = \{+q(2)^D, +p(2)^D, +r(i)^D\}$ ,  $i \in \{3, \dots\}$ , are not constrained as they are not even relevant.  $\square$

We stress that, in order to test whether a relevant (weak) repair  $\mathcal{R}$  is constrained, we need to consider every *subset of occurrences* of non-nullary constants in  $\mathcal{R}$ . For instance, in the case of the repair  $\mathcal{R}_1 = \{q(1), r(1)\}$  from Example 7, the occurrence of the constant 1 in  $q(1)$  is constrained by the presence of  $p(1)$ . Replacing that occurrence of 1 with 3 does not result in a weak repair. However, replacing the occurrence of 1 in  $r$  with 3 gives a weak repair and shows that  $\mathcal{R}_1$  is not constrained.

*Example 8.* Let  $\text{Dom} = \{\perp, 1, 2, \dots\}$  and  $\mathcal{D} = \langle \mathcal{I}, \eta, P \rangle$ , where  $\mathcal{I} = \langle \{q(1, 2), s(1, 2, 3)\}, \emptyset \rangle$ ,  $\eta = \emptyset$  and  $P = \{p(X) \leftarrow q(X, Y), r(X, Y, Z); r(X, Y, Z) \leftarrow s(X, Y, Z), t(X, Y, Z)\}$ . Let us consider the request  $\mathcal{S} = (\{p(1)\}, \emptyset)$ . In this case, our approach yields a unique constrained repair  $\mathcal{R} = \{+t(1, 2, 3)^D\}$ . It recognizes that thanks to  $s(1, 2, 3)$  simply inserting  $t(1, 2, 3)$  guarantees  $r(1, 2, 3)$  to be true and, consequently, ensures the presence of  $p(1)$  in the view. There are other repairs and other relevant repairs, but only the one listed above is constrained.  $\square$

We observe that every (relevant, constrained) weak repair contains a (relevant, constrained) repair.

**Proposition 2.** *Let  $\mathcal{D} = \langle \mathcal{I}, \eta, P \rangle$  be a deductive database and  $\mathcal{S}$  a request. A (relevant, constrained) repair for  $(\mathcal{D}, \mathcal{S})$  exists if and only if a (relevant, constrained) weak repair for  $(\mathcal{D}, \mathcal{S})$  exists.*  $\square$

## 5 Complexity

Finally, we discuss the complexity of decision problems concerning the existence of (weak) repairs of types introduced above. The results we present here have proofs that are non-trivial despite rather strong assumptions we adopted. We present them in the appendix.

We assume that the sets of base and derived predicate symbols, the set of integrity constraints  $\eta$  and the view  $P$  are fixed. The only varying parts in the problems are a database  $\mathcal{I}$  and a request  $\mathcal{S}$ . That is, we consider the *data complexity* setting. Moreover, we assume that  $Dom = \{\perp, 1, 2, \dots\}$ , and take  $=$  and  $\leq$ , both with the standard interpretation on  $\{1, 2, \dots\}$ , as the *only* built-in relations. We restrict integrity constraints to expressions of the form:  $\exists X(\forall Y(A_1 \wedge \dots \wedge A_k \rightarrow B_1 \vee \dots \vee B_m))$ , where  $A_i$  and  $B_i$  are atoms with no occurrences of  $\perp$  constructed of base and built-in predicates, and where every variable occurring in the constraint belongs to  $X \cup Y$ , and occurs in some atom  $A_i$  built of a base predicate.

We start by stating the result on the complexity of deciding the consistency of an indefinite database with integrity constraints. While interesting in its own right, it is also relevant to problems concerning the existence of repairs, as one of the conditions for  $U$  to be a repair is that the database that results from executing  $U$  be consistent.

**Theorem 1.** *The problem to decide whether a database  $\langle \mathcal{I}, \eta \rangle$  has a possible world (is consistent) is NP-complete.*  $\square$

We now turn attention to the problem of checking request satisfaction. Determining the complexity of that task is a key stepping stone to the results on the complexity of deciding whether updates are (weak) repairs that are necessary for our results on the complexity of the existence of (weak) repairs. However, checking request satisfaction turns out to be a challenge even for very simple classes of views. In this paper, we restrict attention to the case when  $P$  is a safe definite (no constraints) acyclic (no recursion) Horn program, although we obtained Proposition 3 in a more general form.

**Proposition 3.** *The problem to decide whether a ground atom  $t$  is true in a deductive database  $\langle \mathcal{I}, \eta, P \rangle$ , where  $\langle \mathcal{I}, \eta \rangle$  is consistent and  $P$  is a safe Horn program, is in the class co-NP.*  $\square$

Next, we consider the problem to decide whether a ground atom  $t$  is false in a deductive database  $\langle \mathcal{I}, \eta, P \rangle$ . We state it separately from the previous one as our present proof of that result requires the assumption of acyclicity.

**Proposition 4.** *The problem to decide whether a ground atom  $t$  is false (ground literal  $\neg t$  is true) in a deductive database  $\langle \mathcal{I}, \eta, P \rangle$ , where  $\langle \mathcal{I}, \eta \rangle$  is consistent and  $P$  is an acyclic Horn program, is in the class co-NP.*  $\square$

With Propositions 3 and 4 in hand, we move on to study the complexity of the problems of the existence of weak repairs. First, we establish an upper bound on the complexity of checking whether an update is a (relevant) weak repair.

**Proposition 5.** *Let  $\mathcal{D} = \langle \mathcal{I}, \eta, P \rangle$ , where  $\eta$  is a set of integrity constraints,  $P$  an acyclic Horn program,  $U$  an update and  $\mathcal{S}$  a request set. The problem of checking whether an update  $U$  is a weak repair (relevant weak repair) for  $(\mathcal{D}, \mathcal{S})$  is in  $\Delta_2^P$ .*  $\square$

With the results above, we can address the question of the complexity of the existence of repairs. The first problem concerns weak repairs and stands apart from others. It turns out, that deciding the existence of a weak repair is NP-complete, which may seem at odds with Proposition 5 (an obvious non-deterministic algorithm guesses an update  $U$  and checks that it is a weak repair apparently performing a “ $\Sigma_2^P$  computation”). However, this low complexity of the problem is simply due to the fact that there are no relevance, constrainedness or minimality constraints imposed on weak repairs. Thus, the question can be reduced to the question whether there is a “small” database  $\mathcal{J}$ , in which the request holds. The corresponding weak repair consists of deleting all elements from  $\mathcal{I}$  and “repopulating” the resulting empty database so that to obtain  $\mathcal{J}$ .

**Theorem 2.** *Let  $\mathcal{D} = \langle \mathcal{I}, \eta, P \rangle$ , where  $\eta$  is a set of integrity constraints, and  $P$  an acyclic Horn program, and let  $\mathcal{S}$  be a request set. The problem of deciding whether there is a weak repair for  $(\mathcal{D}, \mathcal{S})$  is NP-complete.*  $\square$

As noted, the case of the existence of weak repairs is an outlier and deciding the existence of (weak) repairs of other types is much harder (under common assumptions concerning the polynomial hierarchy).

**Theorem 3.** *Let  $\mathcal{D} = \langle \mathcal{I}, \eta, P \rangle$ , where  $\eta$  is a set of integrity constraints, and  $P$  an acyclic Horn program, and let  $\mathcal{S}$  be a request set. The problems of deciding whether there is a relevant weak repair and whether there is a relevant repair for  $(\mathcal{D}, \mathcal{S})$  are  $\Sigma_2^P$ -complete.*  $\square$

The last result concerns constrained (weak) repairs. It provides an upper bound on the complexity of the problem of deciding the existence of constrained repairs. We conjecture that the upper bound is in fact tight but have not been able to prove it. We leave the problem for future work.

**Theorem 4.** *Let  $\mathcal{D} = \langle \mathcal{I}, \eta, P \rangle$ , where  $\eta$  is a set of integrity constraints, and  $P$  an acyclic Horn program, and let  $\mathcal{S}$  be a request set. The problems of deciding whether there is a constrained weak repair and whether there is a constrained repair for  $(\mathcal{D}, \mathcal{S})$  are in  $\Sigma_3^P$ .*  $\square$

## 6 Discussion and conclusion

We presented a declarative framework for view updating and integrity constraint maintenance for indefinite databases. The framework is based on the notion of an indefinite deductive database. In our approach, the indefiniteness appears in the extensional

database and is modeled by a single null value, consistent with the standards of database practice (a condition not followed by earlier works on the view-update problem over indefinite databases). We defined a precise semantics for indefinite deductive databases in terms of possible worlds. We used the framework to formulate and study the view-update problem. Exploiting the concept of minimality of change introduced by an update, we defined several classes of repairs, including relevant and constrained repairs, that translate an update request against a view into an update of the underlying database. Finally, we obtained several complexity results concerning the existence of repairs.

Our paper advances the theory of view updating in three main ways. First, it proposes and studies the setting where extensional databases are indefinite both before and after an update. While *introducing* indefiniteness to narrow down the class of potential repairs was considered before [5], the assumption there was that the initial extensional database was complete. That assumption substantially limits the applicability of the earlier results. Second, our paper proposes a more expressive model of an indefinite extensional database. In our model databases are determined by two sets of facts. The first set of facts specifies what is true and provides an upper bound to what might still be unknown. By CWA, everything else is false. The second set of facts lists exceptions to the “unknown range,” that is, facts that according to the first set might be unknown but are actually false (exceptions). Third, our paper introduces two novel classes of repairs, relevant and constrained, that often substantially narrow down possible ways to fulfill an update request against a view. Relevant repairs do not introduce any new constants and minimize change. Constrained repairs in addition do not involve constants that are in some precise sense “replaceable” and, thus, not grounded in the problem specification.

We already discussed some earlier work on view updating in the introduction as a backdrop to our approach. Expanding on that discussion, we note that the view-update problem is closely related to abduction and is often considered from that perspective. Perhaps the first explicit connection between the two was made by Bry [1] who proposed to use deductive tools as a means for implementing the type of abductive reasoning required in updating views. That idea was pursued by others with modifications that depended on the class and the semantics of the views. For instance, Kakas and Mancarella [9] exploited in their work on view updates the abductive framework by Eshghi and Kowalski [4] and were the first to consider the stable-model semantics for views. Neither of the two works mentioned above was, however, concerned with the case of updates to views over indefinite databases. Console et al. [2] studied the case in which *requests* can involve variables. These variables are replaced by null values and, in this way, null values eventually end up in repaired databases. However, once there, they lose their null value status and are treated just as any other constants. Consequently, no reasoning over null values takes place, in particular, they have no special effect on the notion of minimality. None of the papers discussed studied the complexity of the view-update problem. Instead, the focus was on tailoring resolution-based deductive reasoning tools to handle abduction. Some results on the complexity of abduction for logic programs were obtained by Eiter, Gottlob and Leone [3]. However, again the setting they considered did not assume incompleteness in extensional databases.

Our paper leaves several interesting questions for future work. First, we considered restricted classes of views. That suggests the problem to extend our complexity

results to the full case of Horn programs and, later, stratified ones. Next, we considered a limited class of integrity constraints. Importantly, we disallowed tuple-generating constraints. However, once they are allowed, even a problem of repairing consistency in an extensional database becomes undecidable. A common solution in the database research is to impose syntactic restrictions on the constraints [8]. That suggests considering view-updating in the setting in which only restricted classes of constraints are allowed.

## Acknowledgments

The third author was supported by the NSF grant IIS-0913459. We are grateful to Sergio Greco and Leopoldo Bertossi for helpful discussions.

## References

1. François Bry. Intensional updates: Abduction via deduction. In *Proceedings of ICLP 1990*, pages 561–575, Cambridge, MA, 1990. MIT Press.
2. Luca Console, Maria Luisa Sapino, and Daniele Theseider Dupré. The role of abduction in database view updating. *J. Intell. Inf. Syst.*, 4(3):261–280, 1995.
3. Thomas Eiter, Georg Gottlob, and Nicola Leone. Abduction from logic programs: Semantics and complexity. *Theor. Comput. Sci.*, 189(1-2):129–177, 1997.
4. K. Eshghi and R. A. Kowalski. Abduction compared with negation by failure. In *Proceedings of ICLP 1989*, pages 234–254, Cambridge, MA, 1989. MIT Press.
5. Carles Farré, Ernest Teniente, and Toni Urpí. Handling existential derived predicates in view updating. In *Proceedings of ICLP 2003*, volume 2916 of *LNCS*, pages 148–162. Springer, 2003.
6. Michael Gelfond and Vladimir Lifschitz. The stable model semantics for logic programming. In *Proceedings of ICLP/SLP 1988*, pages 1070–1080, 1988.
7. Michael Gelfond and Vladimir Lifschitz. Classical negation in logic programs and disjunctive databases. *New Generation Comput.*, 9(3/4):365–386, 1991.
8. Sergio Greco, Francesca Spezzano, and Irina Trubitsyna. Stratification criteria and rewriting techniques for checking chase termination. *PVLDB*, 4(11):1158–1168, 2011.
9. A. C. Kakas and P. Mancarella. Database updates through abduction. In *Proceedings of the sixteenth international conference on Very large databases*, pages 650–661, San Francisco, CA, USA, 1990. Morgan Kaufmann Publishers Inc.
10. Leonid Libkin. A semantics-based approach to design of query languages for partial information. In *Semantics in Databases*, volume 1358 of *LNCS*, pages 170–208. Springer, 1995.
11. Enric Mayol and Ernest Teniente. Consistency preserving updates in deductive databases. *IEEE TDKE*, 47(1):61–103, 2003.
12. R. Reiter. On closed world data bases. In H. Gallaire and J. Minker, editors, *Logic and data bases*, pages 55–76. Plenum Press, 1978.
13. Ernest Teniente and Antoni Olivé. Updating knowledge bases while maintaining their consistency. *VLDB J.*, 4(2):193–241, 1995.
14. Stephen Todd. Automatic constraint maintenance and updating defined relations. In *IFIP Congress*, pages 145–148, 1977.
15. Can Türker and Michael Gertz. Semantic integrity support in sql: 1999 and commercial (object-)relational database management systems. *VLDB J.*, 10(4):241–269, 2001.
16. Jeffrey D. Ullman. *Principles of Database and Knowledge-Base Systems, Volume I*. Computer Science Press, 1988.

## Appendix

**Proposition 1.** *Let  $\mathcal{S}$  be a database and  $q$  a fact. Then  $q \in \mathcal{S}_t$  if and only if  $W \models q$ , for every  $W \in \mathcal{W}(\mathcal{S})$ , and  $q \in \mathcal{S}_f$  if and only if  $W \not\models q$ , for every  $W \in \mathcal{W}(\mathcal{S})$ .*

Proof: (1) Let  $q \in \mathcal{S}_t$  and  $W \in \mathcal{W}(\mathcal{S})$ . By Definition 3, we have that  $q \in W^\downarrow$ . Consequently,  $W \models q$ . Conversely, let us suppose that for some atom  $q$ ,  $W \models q$ , for every  $W \in \mathcal{W}(\mathcal{S})$ . We will show that  $q \in \mathcal{S}^\downarrow$ . To this end, let us assume that  $q \notin \mathcal{S}^\downarrow$ .

Let  $q' \in \mathcal{S}$ . We will show that there is a definite fact  $q_d$  such that  $q' \preceq q_d$  and  $q \not\preceq q_d$ . It is evident if  $q'$  and  $q$  are based on a different predicate symbol. Thus, let us assume that  $q' = p(s_1, \dots, s_k)$  and  $q = p(t_1, \dots, t_k)$ , where all  $s_i$  and  $t_i$  are constants from  $Dom$ . If there is  $i$ ,  $1 \leq i \leq k$ , such that  $s_i \neq \perp$ ,  $t_i \neq \perp$  and  $s_i \neq t_i$ , the claim is again evident. Thus, let us assume that for every  $i$ ,  $1 \leq i \leq k$ ,  $s_i = \perp$ , or  $t_i = \perp$  or  $s_i = t_i$ . As  $q \not\preceq q'$ , there is  $i$ ,  $1 \leq i \leq k$ , such that  $s_i = \perp$  and  $t_i = c$  for some  $c \in Dom \setminus \{\perp\}$ . We define  $q_d$  to be any definite fact of the form  $p(u_1, \dots, u_k)$  such that  $q' \preceq q_d$  and  $u_i \neq c$  (such facts exist as  $Dom$  is infinite). Clearly,  $q \not\preceq q_d$  and the claim follows.

We define  $V = \{q_d \mid q' \in \mathcal{S}\}$ . By the definition,  $\mathcal{S} \subseteq V^\downarrow$ . Thus,  $\mathcal{S}_t = \mathcal{S}^\downarrow \subseteq V^\downarrow$ . Furthermore, by the way atom  $q_d$  were selected,  $V \subseteq \mathcal{S}^\uparrow$ . It follows that  $V$  is a possible world for  $\mathcal{S}$ . On the other hand,  $V \not\models q$ , a contradiction.

(2) Let  $q \in \mathcal{S}_f$  and let  $W \in \mathcal{W}(\mathcal{S})$ . If  $W \models q$ , then there is a definite atom  $q_d \in W$  such that  $q \preceq q_d$ . Since  $W \subseteq \mathcal{S}^\uparrow$ , there is an element  $q' \in \mathcal{S}$  such that  $q' \preceq q_d$ . It follows that  $q \approx q'$  and so,  $q \in \mathcal{S}^\approx$ , a contradiction. Thus,  $W \not\models q$ .

Conversely, let us consider an atom  $q$  such that  $W \not\models q$ , for every  $W \in \mathcal{W}(\mathcal{S})$ . Let us suppose that  $q \in \mathcal{S}^\approx$ . Then, there is a definite fact  $q_d \in \mathcal{S}^\uparrow$  such that  $q \preceq q_d$ . Let  $V$  be the set of all definite facts in  $\mathcal{S}^\uparrow$ . Clearly,  $\mathcal{S}_t \subseteq V^\downarrow$  and  $V \subseteq \mathcal{S}^\uparrow$ . Thus,  $V$  is a possible world for  $\mathcal{S}$ . Moreover,  $q_d \in V$ . Thus,  $V \models q$ , a contradiction  $\square$

**Proposition 2.** *Let  $\mathcal{D} = \langle \mathcal{S}, \eta, P \rangle$  be a deductive database and  $\mathcal{S}$  a request. A (relevant, constrained) repair for  $(\mathcal{D}, \mathcal{S})$  exists if and only if a (relevant, constrained) weak repair for  $(\mathcal{D}, \mathcal{S})$  exists.*

Proof: The “only if” part is evident. By the definition, a repair is a weak repair and, similarly, a relevant (constrained, respectively) repair is a relevant (constrained, respectively) weak repair.

For the “if” part, let us consider a weak repair  $U$  for  $(\mathcal{D}, \mathcal{S})$ . If  $V$  is a weak repair for  $(\mathcal{D}, \mathcal{S})$  and  $V \subseteq U$ , then  $NC(\mathcal{D}, \mathcal{S}, V) \subseteq NC(\mathcal{D}, \mathcal{S}, U)$ . It follows that the family of all weak repairs  $V$  such that  $V \subseteq U$  is finite. Thus, it has a minimal element, that is, it contains a repair for  $(\mathcal{D}, \mathcal{S})$ . The argument for relevant (constrained, respectively) weak repairs is similar.  $\square$

*Complexity results.* In order to prove Theorem 1 we need some auxiliary definitions and results. We write  $size(S)$ , for the size of the representation of  $S$ , where  $S$  stands for a deductive database, database, set of integrity constraints, etc.

**Lemma 1.** *If  $\langle \mathcal{S}, \eta \rangle$  is consistent, then it has a possible world  $U$  such that  $|Dom(U)| = O(size(\mathcal{S}) + size(\eta))$ .*



Proof: If  $\eta$  is empty then the claim trivially follows. Let us assume that  $\eta \neq \emptyset$ . Let  $W$  be a possible world for  $\langle \mathcal{S}, \eta \rangle$ . We denote by  $C$  the set of elements from  $Dom_d$  that occur in  $\langle \mathcal{S}, \eta \rangle$ . Let us consider a constraint  $\Phi \in \eta$ . It is of the form

$$\Phi = \exists X \forall Y \varphi,$$

where  $\varphi$  contains no quantifiers (we recall that  $\varphi$  is subject to the syntactic restrictions we imposed on integrity constraints), and  $X \neq \emptyset$ . Let us assume that  $X = \{X_1, \dots, X_j\}$ . Since  $\Phi$  holds in  $W$ , we can select constants  $c_1^\Phi, \dots, c_j^\Phi \in Dom(W)$  so that the constraint

$$\Phi' = \forall Y \varphi(X_1/c_1, \dots, X_j/c_j)$$

holds in  $W$ . Let  $C'$  be the set of all constants of the form  $c_i^\Phi$ , where  $\Phi$  ranges over all constraints in  $\eta$  that contain existential quantifiers. Clearly,  $|C'| = O(\text{size}(\eta))$ .

Let  $\eta'$  be obtained from  $\eta$  by replacing each constraint  $\Phi = \exists X \forall Y \varphi$ , where  $\varphi$  is quantifier-free and  $X \neq \emptyset$ , with the constraint  $\Phi'$ . It is clear that  $\eta'$  consists of universally quantified sentences and that  $W$  is a possible world for  $\langle \mathcal{S}, \eta' \rangle$  (indeed,  $W$  is a possible world for  $\mathcal{S}$  and  $\eta'$  holds in  $W$ ).

Let us now select a possible world  $V$  for  $\langle \mathcal{S}, \eta' \rangle$  so that  $|Dom(V) \setminus (C \cup C')|$  is minimal. As  $W$  is a possible world for  $\langle \mathcal{S}, \eta' \rangle$ ,  $V$  is well defined.

Let  $c \in Dom(V) \setminus (C \cup C')$  and let  $V'$  be the set of all atoms in  $V$  except for those that contain an occurrence of  $c$ . We have that  $\eta'$  holds in  $V'$ . Indeed, let  $\Psi \in \eta'$ , say  $\Psi = \forall Y \varphi$ , where every variable in  $\varphi$  belongs to  $Y$  (and  $\varphi$  is an implication of the form specified above). Let  $\varphi'$  be a ground instance of  $\varphi$  (over  $Dom_d$ ). If  $c$  does not occur in the consequent of  $\varphi'$ , then the truth values of the consequent of  $\varphi'$  in  $V$  and  $V'$ , respectively, are the same, while the truth value of the antecedent can only change from true to false when we move from  $V$  to  $V'$ . Consequently, the truth value of  $\varphi'$  remains true when we move from  $V$  to  $V'$ . If  $c$  occurs in the consequent of  $\varphi'$ , then it also occurs in the base atom in the antecedent of  $\varphi'$ . Thus, the truth value of the antecedent of  $\varphi'$  is false in  $V'$  ( $V'$  contains no atoms with an occurrence of  $c$ ) and so,  $\varphi'$  is true in  $V'$ .

It follows that  $V'$  is not a possible world for  $\mathcal{S}$  (if it were, it would be a possible world for  $\langle \mathcal{S}, \eta' \rangle$  with  $|Dom(V') \setminus (C \cup C')| < |Dom(V) \setminus (C \cup C')|$ , a contradiction with the way  $V$  was chosen). Let  $\mathcal{S} = \langle D, E \rangle$ . Since  $V'$  was obtained from  $V$  by deleting some of its atoms, every atom in  $E$  is false in  $V'$  as it was false in  $V$ . It follows that there is an atom  $A \in D$  such that  $A$  is not true in  $V'$ . Since  $V$  is a possible world of  $\mathcal{S}$ ,  $V \models A$ . Let us assume that  $A$  is definite. Then  $A \in V$  and all constants it contains belong to  $C$ . Thus,  $A \in V'$  and  $V' \models A$ , a contradiction. It follows that  $A$  contains occurrences of  $\perp$ . We assign to  $c$  that atom  $A$ . Let  $k$  be the number of occurrences of  $\perp$  in  $A$ . Then, at most  $k$  constants in  $Dom(V) \setminus (C \cup C')$  can be assigned  $A$ . Indeed, in order for  $A$  to be assigned to a constant  $c \in Dom(V) \setminus (C \cup C')$ , every definite instantiation of  $A$  in  $V$  must contain an occurrence of  $c$  (otherwise, removing all atoms with an occurrence of  $c$  from  $V$  would yield a possible world  $V$  in which  $A$  would be true). Let  $A'$  be any instantiation of  $A$  in  $V$ . If any of the constants assigned  $A$  does not occur in  $A'$ , removing that constant does not result in a possible world in which  $A$  is false. Thus, every constant assigned in  $A$  must occur in  $A'$ , and the claim follows. Let  $a$  be the largest arity of a base predicate. Then, the number of constants in  $|Dom(V) \setminus (C \cup C')|$  is at most  $|D| \times a = O(\text{size}(\mathcal{S}))$  (as  $a$  is fixed, due to the fact that the database schema is fixed). It

follows that  $|Dom(V)| \leq |Dom(V) \setminus (C \cup C')| + |C| + |C'| = O(size(\mathcal{J}) + size(\eta))$ .  $\square$

**Proposition 6.** *If  $\langle \mathcal{J}, \eta \rangle$  is consistent, then it has a possible world  $V$  such that  $size(V) = O(pol(size(\mathcal{J}) + size(\eta)))$ , where  $pol$  is some polynomial independent of  $\langle \mathcal{J}, \eta \rangle$ .*

Proof: Let  $V$  be a possible world for  $\langle \mathcal{J}, \eta \rangle$  constructed in the proof of Lemma 1. By Lemma 1,  $|V| \leq b \times ((size(\mathcal{J}) + size(\eta))^a)$ , where  $b$  is the number of base predicate symbols and  $a$  is the maximum arity of a base predicate (both numbers are fixed and independent of  $\langle \mathcal{J}, \eta \rangle$ ). Without loss of generality, we can assume that all integers in  $V$  that do not appear in  $\langle \mathcal{J}, \eta \rangle$  are not larger than  $Max + O(size(\mathcal{J}) + size(\eta))$ , where  $Max$  is the largest element from  $Dom_d$  occurring in  $\mathcal{J}$  or  $\eta$ , or 0, if  $\langle \mathcal{J}, \eta \rangle$  contains no constants. Indeed, because of our restriction on built-in predicates, the only thing that matters is how those integers relate to each other with respect to  $=$  and  $\leq$ . Thus, all “gaps” among those integers in  $\mathcal{D}(V)$  that are larger than  $Max$  can be eliminated by “shifting” the numbers to the left. It follows that we can assume that the largest integer occurring in  $V$  is of size  $O(size(\mathcal{J}) + size(\eta))$  and so,  $size(V) \leq |V| \times a \times O(size(\mathcal{J}) + size(\eta)) = O((size(\mathcal{J}) + size(\eta))^{a+1})$ . Since  $V'$  is a possible world for  $\langle \mathcal{J}, \eta' \rangle$  it is also a possible world for  $\langle \mathcal{J}, \eta \rangle$ . Thus, the assertion follows.  $\square$

For the argument above it is important that the database schema is fixed. For instance, if  $\mathcal{J} = \langle \{p(\perp, \dots, \perp), q(\perp), r(\perp)\}, \emptyset \rangle$ , where the arity of  $p$  is  $n$  and is not fixed, and  $W = \{q(1), r(2)\} \cup \{p(a_1, \dots, a_n) \mid a_i \in \{1, 2\}\}$ , then no constant can be removed from  $W$ ,  $size(\mathcal{J}) = O(n)$  and  $size(W) = O(n2^n)$ .

**Theorem 1.** *The problem to decide whether a database  $\langle \mathcal{J}, \eta \rangle$  has a possible world (is consistent) is NP-complete.*

Proof: (Membership) Proposition 6 asserts that if  $\langle \mathcal{J}, \eta \rangle$  is consistent, then it has a possible world whose size is polynomial in  $size(\mathcal{J}) + size(\eta)$ . Thus, the following non-deterministic procedure decides the problem: First, the procedure guesses a set of facts  $W$  of size implied by the constructions from Lemma 1 and Proposition 6 (which is polynomial in  $size(\mathcal{J}) + size(\eta)$ ). Next, the procedure checks that  $W$  is a possible world for  $\mathcal{J}$  and that it satisfies all constraints in  $ground(\eta)$ . Since both these tasks can be accomplished in polynomial time, the problem is in the class NP.

(Hardness) The NP-hardness follows by a reduction from the 3-SAT problem to decide whether a quantified boolean formula  $\Phi = \exists X \varphi$  is *true*, where  $X$  is a set of propositional variables and  $\varphi$  is a conjunction of 3-literal clauses over the set of atoms in  $X$ .

Given such a formula  $\Phi$ , we assume that the domain  $Dom$  contains symbols *true* and *false* for representing the truth values, as well as symbols for representing atoms and clauses of the formula  $\varphi$ . With some abuse of the notation, we denote the sets of those symbols with  $X$  and  $Cl$ , respectively. We assume that all those symbols are distinct and different from the domain elements 1, 2 and 3, which we use to represent positions of literals in clauses (for each clause we fix the order).

Further, we assume that we have base predicate symbols *val*, *sat*, and *occur*. The predicates *val* and *sat* are binary. We use *val* to specify the truth value of atoms in  $X$  and *sat* to specify the truth value of clauses in  $Cl$  (given the assignment of truth values to atoms as determined by *val*). Finally, the predicate *occur* is 4-ary. If an atom  $x$  occurs in

a clause  $c \in Cl$  positively (negatively, respectively) in the position  $p \in \{1, 2, 3\}$ , we represent that by an atom  $occur(c, p, x, true)$  ( $occur(c, p, x, false)$ , respectively). Clearly, all atoms of the form  $occur(c, p, x, v)$  uniquely determine the formula  $\phi$ . We define  $\mathcal{J} = \langle D, \emptyset \rangle$  where  $D$  consists of:

1.  $val(x, \perp)$ , for each  $x \in X$
2.  $sat(c, \perp)$ , for each  $c \in Cl$
3.  $occur(c, p, x, true)$ , if  $x$  occurs positively in  $c$  in a position  $p \in \{1, 2, 3\}$
4.  $occur(c, p, x, false)$ , if  $x$  occurs negatively in  $c$  in a position  $p \in \{1, 2, 3\}$ .

Next, we define  $\eta$  to consist of the following constraints (we note all of them satisfy the syntactic “safety” restriction we imposed):

1.  $\forall A, V (val(A, V) \rightarrow V = true \vee V = false)$
2.  $\forall A (val(A, true) \wedge val(A, false) \rightarrow \perp)$
3.  $\forall C, V (sat(C, V) \rightarrow V = true \vee V = false)$
4.  $\forall C (sat(C, true) \wedge sat(C, false) \rightarrow \perp)$
5.  $\forall C, P, A, V (sat(C, false) \wedge occur(C, P, A, V) \wedge val(A, V) \rightarrow \perp)$
6.  $\forall C, A, A', A'', V, V', V''$   
 $(sat(C, true) \wedge occur(C, 1, A, V) \wedge \neg val(A, V)$   
 $\wedge occur(C, 2, A', V') \wedge \neg val(A', V')$   
 $\wedge occur(C, 3, A'', V'') \wedge \neg val(A'', V'') \rightarrow \perp)$
7.  $\forall C (sat(C, false) \rightarrow \perp)$ .

One can prove that  $\Phi$  is true if and only if  $\langle \mathcal{J}, \eta \rangle$  is consistent.  $\square$

**Proposition 3.** *The problem to decide whether a ground atom  $t$  is true in a deductive database  $\langle \mathcal{J}, \eta, P \rangle$ , where  $\langle \mathcal{J}, \eta \rangle$  is consistent and  $P$  is a safe Horn program, is in the class co-NP.*

Proof: It suffices to show that the problem to decide whether  $t$  is *not* true in a deductive database  $\langle \mathcal{J}, \eta, P \rangle$  is in the class NP. That problem consists of deciding whether there is a possible world  $W$  of  $\langle \mathcal{J}, \eta \rangle$  such that  $LM(P \cup W) \not\models t$ . If such a possible world  $W$  exists, the construction from the proof of Proposition 6 shows that there is a possible world  $V \subseteq W$  of the size polynomial in the size of  $\mathcal{J}$ . Since  $V \subseteq W$ , we have  $LM(P \cup V) \not\models t$ . Thus, to decide that  $t$  is not true in  $\langle \mathcal{J}, \eta, P \rangle$ , one guesses  $V$ , a polynomial-size possible world for  $\langle \mathcal{J}, \eta \rangle$ , and then verifies that it is a possible world for  $\langle \mathcal{J}, \eta \rangle$  and that  $t$  does not hold in it. These verification tasks are polynomial in the size of  $\mathcal{J}$  and so, the problem to decide whether  $t$  is not true in a deductive database  $\langle \mathcal{J}, \eta, P \rangle$  is in the class NP, as required.  $\square$

**Proposition 4.** *The problem to decide whether a ground atom  $t$  is false (ground literal  $\neg t$  is true) in a deductive database  $\langle \mathcal{J}, \eta, P \rangle$ , where  $\langle \mathcal{J}, \eta \rangle$  is consistent and  $P$  is an acyclic Horn program, is in the class co-NP.*

Proof: To prove the assertion, we show that the complementary problem to decide whether  $t$  is *not* false in  $\langle \mathcal{J}, \eta, P \rangle$  is in NP. To this end, it suffices to show that if  $t$  is not false in  $\langle \mathcal{J}, \eta, P \rangle$  then there is a possible world  $V$  of  $\langle \mathcal{J}, \eta \rangle$  such that the size of  $V$  is polynomial in the size of  $\mathcal{J}$  and  $P \cup V \models t$ .

Clearly, if  $t$  is not false in  $\langle \mathcal{S}, \eta, P \rangle$ , then there is a possible world  $W$  of  $\langle \mathcal{S}, \eta \rangle$  such that  $P \cup W \models t$ . Let  $T$  be the tree of a shortest resolution proof of  $t$  from  $\text{ground}(P \cup W)$  and  $W_T$  the subset of atoms in  $W$  occurring in  $T$ . Let  $V$  be any subset of  $W$  such that

1.  $(W \cap \mathcal{S}_t) \cup W_T \subseteq V$
2.  $V$  is minimal with respect to that requirement.

In other words, for  $V$  we take any minimal subset of  $W$  that still “explains” the truth of every true atom of  $\mathcal{S}$  and contains all atoms  $W_T$ . We stress that there are many of these sets. We have the following properties:

1.  $LM(P \cup V) \models t$  (holds, as  $W_T \subseteq V$ )
2. The size of  $V$  is polynomial in the size of  $\mathcal{S}$  (follows by the fact that  $P$  is not recursive and  $T$  is a tree of a *shortest* proof, which implies that  $|W_T|$  is bounded by a constant independent of the size of  $\mathcal{S}$ ).

As  $V$  is a subset of  $W$ , it may happen that some integrity constraints are violated. Specifically, it may now be that the body of an integrity constraint is *true* with respect to  $W$ , remains *true* with respect of  $V$ , and its head, which is *true* in  $W$ , becomes *false* in  $V$ . In order to “fix” that constraints we reinsert into  $V$  atoms from  $W$  that make the consequent of the constraint true or, to be precise, a minimal set of such atoms. In this way, we obtain a new set, which we substitute for  $V$ . We continue this process as long as there are violated integrity constraints. The key point is that this process can only reintroduce a polynomial number of atoms because the integrity constraints are safe (no existential quantification in the consequents) and the constants involved in the process are only those occurring in  $\mathcal{S}$  and in  $T$ . Therefore, when the process terminates,  $V$  is a possible world of  $\langle \mathcal{S}, \eta \rangle$ , its size is polynomial in the size of  $\mathcal{S}$  and  $LM(P \cup V) \models t$ .  $\square$

We believe that the result above can be extended to the case of (non-acyclic) Horn programs. However, a proof of such a result would have to follow a different approach. The reason is that given a possible world  $W$  for  $\langle \mathcal{S}, \eta \rangle$  such that  $LM(P \cup W) \models t$ , it is not always possible to find a possible-world  $V \subseteq W$  of size polynomial in the size of  $\langle \mathcal{S}, \eta \rangle$  such that  $LM(P \cup V) \models t$ . Let  $\mathcal{S} = \langle \{start(1), edge(\perp, \perp)\}, \emptyset \rangle$ ,  $\eta = \emptyset$  and  $P$  consist of:

$$\begin{aligned} r(X) &\leftarrow start(X) \\ r(X) &\leftarrow r(Y), edge(Y, X), r(X). \end{aligned}$$

Let us define  $t = r(0)$  and  $W = \{start(1), edge(1, 2), edge(2, 3), \dots, edge(n-1, n), edge(n, 0)\}$ . Then  $W$  is a possible world for  $\langle \mathcal{S}, \eta \rangle$  and  $LM(P \cup W) \models t$ . But  $W$  contains no fixed-size possible world  $V$  for  $\langle \mathcal{S}, \eta \rangle$  such that  $LM(P \cup V) \models t$ . Such a fixed-size possible world exists but has to be constructed independently of any particular possible world.

**Proposition 5.** Let  $\mathcal{D} = \langle \mathcal{S}, \eta, P \rangle$ , where  $\eta$  is a set of integrity constraints,  $P$  an acyclic Horn program,  $U$  an update and  $\mathcal{S}$  a request set. The problem of checking whether an update  $U$  is a weak repair for  $(\mathcal{D}, \mathcal{S})$  is in  $\Delta_2^P$ .

Proof: Let  $\mathcal{D}' = \langle \mathcal{S}', \eta, P \rangle$ , where  $\mathcal{S}' = \mathcal{S} \circ U$ . In order to test whether  $U$  is a weak repair we need to check whether  $\mathcal{D}'$  is consistent and  $\mathcal{D}' \models \mathcal{S}$ . Clearly,  $\mathcal{D}'$  can be constructed in polynomial time. The first problem is NP-complete and so, can be solved by

a call to an NP oracle (Theorem 1). The second one can be solved with a polynomial number of calls to coNP oracles (cf. Propositions 3 and 4).  $\square$

**Theorem 2.** *Let  $\mathcal{D} = \langle \mathcal{I}, \eta, P \rangle$ , where  $\eta$  is a set of integrity constraints, and  $P$  an acyclic Horn program, and let  $\mathcal{S}$  be a request set. The problem of deciding whether there is a weak repair for  $(\mathcal{D}, \mathcal{S})$  is NP-complete.*

*Proof:* (Membership) Let us assume that there is a weak repair for  $(\mathcal{D}, \mathcal{S})$ . It follows that there is an indefinite database  $\mathcal{J}$  that satisfies the request. Let  $W$  be a possible world of  $\mathcal{J}$ . Let  $a$  be an atom required to be true. Since  $\mathcal{J}$  satisfies the request, there is a definite atom  $b$  such that  $a \preceq b$  and  $b$  has a proof from  $P \cup W$ . Without loss of generality we may assume that the proof is minimal (no formula can be eliminated). It follows that the cardinality of that proof does not depend on  $\mathcal{I}$  nor on  $\mathcal{J}$  nor on  $\mathcal{S}$ . It is so because  $P$  is acyclic and fixed (it is independent of the database component and of the request). Let  $W_a$  be the set of atoms from  $W$  occurring in that proof,  $C_a$  the set of all constants occurring in that proof and let  $W'$  and  $C$  be the unions of all sets  $W_a$  and  $C_a$ , respectively, over all atoms  $a$  required to be true. Finally, we define  $W''$  to be the set of all atoms in  $W$  that contain only occurrences of constants in  $C$ . It is clear that  $W' \subseteq W'' \subseteq W$  and that the size of  $W''$  is polynomial in the size of  $\mathcal{S}$ . We will show that the request  $\mathcal{S}$  holds in  $W''$  and that  $W''$  satisfies the integrity constraints.

First, since  $W' \subseteq W''$ ,  $P \cup W'' \models a$ , for every atom  $a$  requested to be true. Second, since for every atom  $b$  requested to be false,  $P \cup W \not\models b$ ,  $P \cup W'' \not\models b$ . Both properties follow from the fact that  $P$  is a Horn program and  $W' \subseteq W'' \subseteq W$ .

Next, we note that  $W''$  is obtained from  $W$  by removing from  $W$  all atoms that contain at least one constant not in  $C$ . Since  $W$  satisfies all the integrity constraints and since for every ground instance of an integrity constraint and for every constant that occurs in the consequent there is a base predicate in the antecedent with an occurrence of that constant,  $W''$  satisfies all integrity constraints, as well.

To recap, we proved that if there is a weak repair for  $(\mathcal{D}, \mathcal{S})$ , then there is a set of definite atoms  $W''$  built of base predicates such that the size of  $W''$  is polynomial in the size of  $\mathcal{S}$ ,  $W''$  satisfies all the integrity constraints and  $P \cup W'' \models \mathcal{S}$ . Conversely, if such a set  $W''$  exists, then there is a weak repair for  $(\mathcal{D}, \mathcal{S})$ . Indeed, such a repair might be constructed of all update atoms needed to transform  $\mathcal{J}$  into an empty database and all insert atoms converting that empty database into  $W''$ . Since  $W''$  is definite,  $W''$  is the only possible world of  $\langle W'', \eta, P \rangle$  and by the properties of  $W''$ ,  $\langle W'', \eta, P \rangle$  satisfies the request.

It follows that the following algorithm decides whether there is a weak repair for  $(\mathcal{D}, \mathcal{S})$ : guess a polynomial-size definite database  $W''$  and verify that it satisfies the integrity constraints and the request. Both checks can be performed in polynomial time. Thus, the problem is in the class NP.

(Hardness) We prove the claim by a reduction from the 3-SAT problem. That problem consists of deciding whether a formula  $\Phi = \exists X \varphi$  is *true*, where  $X$  is a set of propositional atoms and  $\varphi$  is a CNF formula over  $X$  with each clause being a disjunction of 3 literals. We denote the set of clauses in  $\varphi$  by  $Cl$ .

Given such a formula  $\Phi$ , we assume that the domain  $Dom$  contains symbols *true* and *false* to represent the truth values, and symbols to represent atoms and clauses of

$\phi$ . With some abuse of notation we will denote those two sets of symbols by  $X$  and  $Cl$ , respectively. We also assume these symbols are different from 1, 2 and 3, which we will use to denote positions of literals in clauses (we fix the order of literals in each clause). Next, we assume that we have base predicates  $val$  and  $occur$ . The predicate  $val$  is binary, and the predicate  $occur$  is 4-ary. We use  $val$  to define the truth value of a propositional atom in  $X$ . If an atom  $x$  occurs in a clause  $c$  positively (negatively, respectively) in a position  $p \in \{1, 2, 3\}$ , we represent that by an atom  $occur(c, p, x, true)$  ( $occur(c, p, x, false)$ , respectively). As derived predicates, we use predicates  $sat$  and  $occur'$ . The predicate  $sat$  is unary; we use it to represent the truth value of clauses in  $\phi$ . The predicate  $occur'$  has the same arity and plays the same role as its non-primed counterpart. We define  $\mathcal{S} = \langle \emptyset, \emptyset \rangle$  and  $\eta$  to consist of the following constraints:

1.  $\forall C, P, A, A' V (occur(C, P, A, V) \wedge occur(C, P, A', V) \rightarrow A = A')$
2.  $\forall C, P, A, V (occur(C, P, A, V) \rightarrow P = 1 \vee P = 2 \vee P = 3)$
3.  $\forall C, P, A, V (occur(C, P, A, V) \rightarrow V = true \vee V = false)$
4.  $\forall C, P, A (occur(C, P, A, true) \wedge occur(C, P, A, false) \rightarrow \perp)$
5.  $\forall A, V (val(A, V) \rightarrow V = true \vee V = false)$
6.  $\forall A (val(A, true) \wedge val(A, false) \rightarrow \perp)$ .

Next, we define  $P$  to consist of the following rules:

1.  $occur'(C, P, W, V) \leftarrow occur(C, P, W, V)$
2.  $sat(C) \leftarrow occur(C, P, W, V), val(W, V)$

Finally, we define the request set  $\mathcal{S}$  to consist of the following literals:

1.  $occur'(c, p, x, true)$ , for every clause  $c$  in  $\phi$ ,  $x \in X$  and  $p \in \{1, 2, 3\}$  such that  $x$  occurs positively in  $c$  in position  $p$
2.  $occur'(c, p, x, false)$ , for every clause  $c$  in  $\phi$ ,  $x \in X$  and  $p \in \{1, 2, 3\}$  such that  $x$  occurs negatively in  $c$  in position  $p$
3.  $sat(c)$ , for every  $c \in Cl$ .

One can prove that  $\Phi$  is *true* if and only if a weak repair for  $(\mathcal{D}, \mathcal{S})$  exists.  $\square$

**Theorem 3.** Let  $\mathcal{D} = \langle \mathcal{S}, \eta, P \rangle$ , where  $\eta$  is a set of integrity constraints, and  $P$  an acyclic Horn program, and let  $\mathcal{S}$  be a request set. The problems of deciding whether there is a relevant weak repair and whether there is a relevant repair for  $(\mathcal{D}, \mathcal{S})$  are  $\Sigma_2^P$ -complete.

*Proof:* Firstly observe that there are only finitely many relevant weak repairs. Therefore, a  $\sqsubseteq$ -minimal relevant weak repair exists if and only if a relevant weak repair exists. In the following we prove that the problem of deciding whether there is a relevant weak repair for  $(\mathcal{D}, \mathcal{S})$  is  $\Sigma_2^P$ -complete.

(Membership) Each relevant weak repair for  $(\mathcal{D}, \mathcal{S})$  consists of insert and delete actions involving predicate symbols and constants occurring in  $\mathcal{D}$  and  $\mathcal{S}$  and the additional constant  $\perp$  representing the null value. Therefore, every relevant weak repair has size polynomial in the size of  $(\mathcal{D}, \mathcal{S})$ . Once we guess a relevant update  $U$ , we have to test whether the updated database  $\mathcal{D} \circ U$  is consistent, which can be decided by an NP



oracle (Theorem 1), and whether the request is satisfied, which can be decided by polynomially many calls to co-NP oracles (Propositions 3 and 4). Thus, the membership of the problem in the class  $\Sigma_2^P$  follows.

(Hardness) In order to show  $\Sigma_2^P$ -hardness, we present a reduction from the problem to decide whether a quantified boolean formula  $\Phi = \exists X \forall Y \varphi$  is true, where  $X$  and  $Y$  are disjoint sets of propositional atoms and  $\varphi$  is a 3-DNF formula over the set of atoms  $X \cup Y$ . It is indeed sufficient, as that problem is known to be  $\Sigma_2^P$ -complete.

Given such a formula  $\Phi$ , we denote by  $D_\varphi$  the set of (identifiers of) disjuncts in  $\varphi$ . Clearly, each  $d \in D_\varphi$  is a conjunction of literals. We denote by  $\varphi_d$  the set of atoms that occur in  $d$ . We assume that the domain  $Dom$  contains symbols *true* and *false*, as well as symbols to represent atoms in  $X$  and in  $Y$ , and disjuncts of  $\varphi$  (elements of  $D_\varphi$ ). With some abuse of notation we will denote those sets of symbols by  $X$ ,  $Y$  and  $D_\varphi$ , respectively. We assume that all these symbols are different from three additional domain symbols, 1, 2 and 3. We denote the set of all constants mentioned above by  $Dom_r$ . That is, we set  $Dom_r = X \cup Y \cup D_\varphi \cup \{1, 2, 3\} \cup \{true, false\}$ .

We assume base predicates  $inX$ ,  $inY$ ,  $disj$ ,  $occur$ ,  $inX_c$ ,  $inY_c$ ,  $disj_c$ ,  $occur_c$ ,  $val_X$ ,  $val_Y$  and  $assign$ . The predicates  $inX$ ,  $inY$  are unary and their extensions give the sets of atoms  $X$  and  $Y$ . The predicate  $disj$  is also unary; its extension specifies the disjuncts of the formula  $\varphi$ . The predicate  $occur$  is 4-ary. If an atom  $a$  occurs in a disjunct  $d$  positively (negatively, respectively) in a position  $p \in \{1, 2, 3\}$ , we represent that by an atom  $occur(d, p, x, true)$  ( $occur(d, p, x, false)$ , respectively). The predicates with the subscript  $c$  have the same arities as their subscript-free counterparts and are intended to represent their complements with respect to  $Dom_r$ . The predicates  $val_X$ ,  $val_Y$  and  $assign$  are binary. We use  $val_X$  and  $val_Y$  to specify the truth value of atoms in  $X$  and  $Y$ , respectively. The predicate  $assign$  will associate some new constants (not in  $Dom_r$ ) with *true* and some other new constants with *false* (the need for that will become evident from the proof below).

We also assume derived predicates  $q'$  for  $q \in \{inX, inX_c, inY, inY_c, disj, disj_c, occur, occur_c\}$ , each of the same arity as its non-primed counterpart, as well as three additional ones  $val$ ,  $sat$  and  $satisfied$ . The predicates  $q'$ ,  $q \in \{inX, inX_c, inY, inY_c, disj, disj_c, occur, occur_c\}$ , allow us to formulate a request so that no relevant weak repair can change the extensions of the corresponding predicates in the database. The predicate  $val$  is unary and allows us to formulate a request that ensures that every possible world of the database resulting from a relevant weak repair assigns a truth value to every atom in  $X \cup Y$ . The  $sat$  is unary and  $satisfied$  is propositional (0-ary). We use them to represent the satisfaction of individual disjuncts in  $\varphi$  and of  $\varphi$  itself.

We now define the database  $\mathcal{J} = \langle D, \emptyset \rangle$  where  $D$  consists of the following atoms:

1.  $inX(a)$ , for every atom  $a \in X$
2.  $inY(a)$ , for every atom  $a \in Y$
3.  $disj(d)$ , for every disjunct  $d \in D_\varphi$
4.  $occur(d, a, p, true)$  for every atom  $a \in X \cup Y$  that occurs non-negated in position  $p$  in the conjunct  $d$
5.  $occur(d, a, p, false)$ , for every atom  $a \in X \cup Y$  that occurs negated in position  $p$  in the conjunct  $d$

6. Finally, we include in  $D$  atoms  $inX_c(a)$ ,  $inY_c(a)$ ,  $disj_c(d)$ ,  $occur_c(d, a, p, v)$  so that the extensions of  $inX_c$ ,  $inY_c$ ,  $disj_c$  and  $occur_c$  are the complements with respect to  $Dom_r$  ( $Dom_r^4$ , in the last case) of the extensions of the corresponding predicates as specified in (1)-(5).

We define  $\eta$  to consist of the constraints:

1.  $\forall A, V (val_X(A, V) \rightarrow V = true \vee V = false)$
2.  $\forall A (val_X(A, true) \wedge val_X(A, false) \rightarrow \perp)$
3.  $\forall A, V (val_X(A, V) \rightarrow inX(A))$
4.  $\forall B, V (val_Y(B, V) \rightarrow inY(B))$
5.  $\forall B (val_Y(B, true) \rightarrow \perp)$
6.  $\forall B (val_Y(B, false) \rightarrow \perp)$
7.  $\forall B (val_Y(B, 1) \rightarrow \perp)$
8.  $\forall B (val_Y(B, 2) \rightarrow \perp)$
9.  $\forall B (val_Y(B, 3) \rightarrow \perp)$
10.  $\forall B, A (inX(A) \wedge val_Y(B, A) \rightarrow \perp)$
11.  $\forall B, A (inY(A) \wedge val_Y(B, A) \rightarrow \perp)$
12.  $\forall B, D (disj(D) \wedge val_Y(B, D) \rightarrow \perp)$
13.  $\forall B, V, V' (val_Y(B, V) \wedge val_Y(B, V') \rightarrow V = V')$
14.  $\forall B, V (val_Y(B, V) \rightarrow assign(V, true) \vee assign(V, false))$
15.  $\forall B, V (val_Y(B, V) \wedge assign(V, true) \wedge assign(V, false) \rightarrow \perp).$

The view  $P$  consists of the following rules:

1.  $q'(\mathbf{W}) \leftarrow q(\mathbf{W})$ , for every base predicate  $q \in \{inX, inX_c, inY, inY_c, disj, disj_c, occur, occur_c\}$  ( $\mathbf{W}$  stands for a vector of different variable symbols of the appropriate arity)
2.  $val(A) \leftarrow inX(A), val_X(A, V)$
3.  $val(A) \leftarrow inY(A), val_Y(A, V)$
4.  $sat(D, P) \leftarrow occur(D, P, A, V), inX(A), val_X(A, V)$
5.  $sat(D, P) \leftarrow occur(D, P, A, V), inY(A), val_Y(A, V'), assign(V', V)$
6.  $satisfied \leftarrow sat(D, 1), sat(D, 2), sat(D, 3).$

Finally, we define the request set  $\mathcal{S}_\phi$  to consist of the facts:

1.  $q'(\mathbf{a})$ , for every atom of the form  $q(\mathbf{a})$  in  $\mathcal{S}$  ( $q \in \{inX, inX_c, inY, inY_c, disj, disj_c, occur, occur_c\}$ ).
2.  $val(A)$ , for every  $A \in X \cup Y$
3.  $satisfied$ .

We note that the part (1) of the request guarantees that no matter what relevant weak repair we consider, the extensions of base predicates  $inX$ ,  $inX_c$ ,  $inY$ ,  $inY_c$ ,  $disj$ ,  $disj_c$ ,  $occur$  and  $occur_c$  will be the same before and after the repair is applied. Atoms given in the specification (2) of the request guarantee that applying any weak repair results in a database that contains  $val_X(\perp, \perp)$  or, for every  $a \in X$ ,  $val_X(a, true)$ ,  $val_X(a, false)$  or  $val_X(a, \perp)$ . Since no constants present in the original database or in the request can be used as  $\alpha$  in atoms of the form  $val_Y(y, \alpha)$ , it also follows that every relevant weak

repair inserts into the database atom  $val_Y(\perp, \perp)$  or, for every  $b \in Y$ , atoms  $val_Y(y, \perp)$ . Due to the integrity constraint (14), it also inserts  $assign(\perp, \perp)$  or  $assign(\perp, true)$  and  $assign(\perp, false)$ .

We now prove that there exists a relevant weak repair for the database  $(\mathcal{D}(\Phi), \mathcal{S}(\Phi))$  if and only if  $\Phi = \exists X \forall Y \varphi$  is true. Let  $W$  be a set of propositional atoms. An *assignment*  $\Theta_W$  of truth values to atoms in  $W$  determines a function  $\Theta_W(\cdot)$ , that takes as its argument a formula  $\varphi$  and returns the formula obtained by replacing in  $\varphi$  each  $w \in W$  with a truth value (*true* or *false*). Given two assignments  $\Theta_X(\cdot)$  and  $\Theta_Y(\cdot)$ , we define the function  $\Theta_{X,Y}(\cdot) = \Theta_X(\Theta_Y(\cdot))$ .

( $\Rightarrow$ ) Let  $U$  be a relevant weak repair for the database  $(\mathcal{D}(\Phi), \mathcal{S}(\Phi))$ . Our goal is to derive from  $U$  an assignment  $\Theta_X$  such that for each assignment  $\Theta_Y$  the truth value of  $\Theta_{X,Y}(\varphi)$  is *true*.

Let  $\mathcal{S}' = \mathcal{S} \circ U$  and  $x \in X$ . By our discussion above, we have the following cases:

1.  $\mathcal{S}'$  contains an atom of the form  $val_X(x, true)$  ( $val_X(x, false)$ , respectively). In this case we set  $\Theta_X(x) = true$  ( $\Theta_X(x) = false$ , respectively). We note that  $\mathcal{S}'$  cannot contain both  $val_X(x, true)$  and  $val_X(x, false)$  as that would violate one of the integrity constraints.
2.  $\mathcal{S}'$  does not contain an atom of the form  $val_X(x, true)$  or  $val_X(x, false)$ . Then,  $\mathcal{S}'$  contains  $val_X(x, \perp)$  or  $val_X(\perp, \perp)$ . In this case, the truth value of  $X$  can be arbitrary and we set  $\Theta_X(x) = true$ .

We will now show that for each assignment  $\Theta_Y$  of truth values to atoms in  $Y$ ,  $\Theta_{X,Y}(\varphi)$  is *true*. By our discussion above, since  $U$  is a relevant weak repair,  $\mathcal{S}'$  contains the atom  $val_Y(\perp, \perp)$  or, for every  $y \in Y$ , the atom  $val_Y(y, \perp)$ , and it also contains the atom  $assign(\perp, \perp)$  or, the atoms  $assign(\perp, true)$  and  $assign(\perp, false)$ .

Let  $\xi_t$  and  $\xi_f$  be domain elements not in  $Dom_r$ . We define

$$\begin{aligned} W = & \mathcal{S} \cup \{val_X(x, \Theta_X(x)) \mid x \in X\} \\ & \cup \{val_Y(y, \xi_t) \mid y \in Y, \Theta_Y(y) = true\} \\ & \cup \{assign(\xi_t, true), assign(\xi_f, false)\}. \end{aligned}$$

It is straightforward to check that  $W$  is a possible world of  $\langle \mathcal{S}', \eta \rangle$ . Thus, *satisfied* holds in  $LM(P \cup W)$ . Consequently, there is a disjunct  $d \in D_\varphi$  such that  $sat(d)$  holds in  $LM(P \cup W)$ . It follows that  $\Theta_{X,Y}(\varphi) = true$ .

( $\Leftarrow$ ) Now, let us suppose that there is an assignment  $\Theta_X$  of truth value to atoms in  $X$  such that for each assignment  $\Theta_Y$  of truth values to atoms in  $Y$ , the truth value of  $\Theta_{X,Y}(\varphi)$  is true. We define  $U = \{+val_X(x, \Theta_X(x)), +val_Y(\perp, \perp), assign(\perp, \perp)\}$ . The revised database  $\mathcal{S}'$  satisfies  $\mathcal{S}' = \mathcal{S} \cup \{+val_X(x, \Theta_X(x)), +val_Y(\perp, \perp), assign(\perp, \perp)\}$ . It is easy to check that  $\langle \mathcal{S}', \eta \rangle$  is consistent (we recall that  $\langle \mathcal{S}, \eta \rangle$  is consistent).

It is clear that the parts (1) and (2) of the request hold after the repair. We will prove that *satisfied* holds after the repair. Let us consider an arbitrary possible world  $W$  of  $\langle \mathcal{S}', \eta \rangle$ . The integrity constraints imply that for every  $y \in Y$  there is exactly one atom  $val_Y(y, \zeta_y)$  and exactly one atom  $assign(\zeta_y, v)$ , where  $v = true$  or  $v = false$ . We define  $\Theta_Y$  so that for every  $y \in Y$ ,  $\Theta_Y(y) = true$  if  $assign(\zeta_y, true) \in W$ , and  $\Theta_Y(y) = false$ , otherwise. Since  $\Theta_{X,Y}(\varphi)$  holds, there is a disjunct  $d \in D_\varphi$  such that  $\Theta_{X,Y}(d)$  holds.

Thus,  $\text{sat}(d)$  holds in  $LM(P \cup W)$  and so,  $\text{satisfied}$  holds in  $LM(P \cup W)$ , too.  $\square$

In order to prove Theorem 4 we need the following proposition.

**Proposition 7.** *Let  $\mathcal{D} = \langle \mathcal{I}, \eta, P \rangle$ , where  $\eta$  is a set of integrity constraints,  $P$  an acyclic Horn program,  $U$  an update and  $\mathcal{S}$  a request set. The problem of deciding whether an update  $U$  is a constrained weak repair for  $(\mathcal{D}, \mathcal{S})$  is in  $\Pi_2^P$ .*

Proof: By the definition to decide whether  $U$  is not a constrained weak repair we need to guess a non-nullary constant  $a$  and a subset  $Q$  of its occurrences and show that:

1.  $U$  is a weak repair for  $(\mathcal{D}, \mathcal{S})$ , and
2. the update  $U'$  resulting from replacing occurrences of  $a$  in  $Q$  is a weak repair for  $(\mathcal{D}, \mathcal{S})$ .

Both tasks are in  $\Delta_2^P$  (Proposition 5). Therefore, the problem of deciding whether  $U$  is not a constrained weak repair is in  $\Sigma_2^P$  and the problem of deciding whether  $U$  is a constrained weak repair is in  $\Pi_2^P$ .  $\square$

**Theorem 4.** *Let  $\mathcal{D} = \langle \mathcal{I}, \eta, P \rangle$ , where  $\eta$  is a set of integrity constraints, and  $P$  an acyclic Horn program, and let  $\mathcal{S}$  be a request set. The problems of deciding whether there is a constrained weak repair and whether there is a constrained repair for  $(\mathcal{D}, \mathcal{S})$  are in  $\Sigma_3^P$ .*

Proof: Firstly observe that there are only finitely many constrained weak repairs. Therefore, a  $\sqsubseteq$ -minimal constrained weak repair exists if and only if a constrained weak repair exists. In the following we prove that the problem of deciding whether there is a constrained weak repair for  $(\mathcal{D}, \mathcal{S})$  is  $\Sigma_3^P$ -complete.

Each constrained weak repair is a relevant weak repair. Therefore, its size is polynomial in the size of  $\mathcal{D}$ . Once we guess a relevant update  $U$ , we have to test if it is a constrained weak repair. We know, by Proposition 7, that the second problem is in  $\Pi_2^P$ . Therefore, the problem in the class  $NP^{\Pi_2^P} = \Sigma_3^P$ .  $\square$